arXiv:1602.00159v3 [q-fin.EC] 6 Jun 2016

Empirical Methods for Dynamic Power Law Distributions in the Social Sciences

Ricardo T. Fernholz*

Claremont McKenna College

June 7, 2016

Abstract

This paper introduces nonparametric econometric methods that characterize general power law distributions under basic stability conditions. These methods extend the literature on power laws in the social sciences in several directions. First, we show that any stationary distribution in a random growth setting is shaped entirely by two factors—the idiosyncratic volatilities and reversion rates (a measure of cross-sectional mean reversion) for different ranks in the distribution. This result is valid regardless of how growth rates and volatilities vary across different economic agents, and hence applies to Gibrat's law and its extensions. Second, we present techniques to estimate these two factors using panel data. Third, we show how our results offer a structural explanation for a generalized size effect in which higher-ranked processes grow more slowly than lower-ranked processes on average. Finally, we employ our empirical methods using data on commodity prices and show that our techniques accurately describe the empirical distribution of relative commodity prices. We also show the existence of a generalized "size" effect for commodities, as predicted by our econometric theory.

JEL Codes: C10, C14

Keywords: power laws, Pareto distribution, Gibrat's Law, nonparametric methods, size effect, commodity prices

*Robert Day School of Economics and Finance, Claremont McKenna College, 500 E. Ninth St., Claremont, CA 91711, rfernholz@cmc.edu.

1 Introduction

Power laws are ubiquitous in economics, finance, and the social sciences more broadly. They are found across many different phenomena, ranging from the distribution of income and wealth (Atkinson et al., 2011; Piketty, 2014) to the city size distribution (Gabaix, 1999) to the distribution of assets of financial intermediaries (Janicki and Prescott, 2006; Fernholz and Koch, 2016). Although a number of potential mechanisms explaining the appearance of power laws have been proposed (Newman, 2006; Gabaix, 2009), one of the most broad and influential involves random growth processes.

A large literature in economics, both theoretical and empirical, models different power laws and Pareto distributions as the result of random growth processes that are stabilized by the presence of some friction (Champernowne, 1953; Gabaix, 1999; Luttmer, 2007; Benhabib et al., 2011). In this paper, we present rank-based, nonparametric methods that allow for the characterization of general power law distributions in any continuous random growth setting. These techniques, which are well-established and the subject of active research in statistics and mathematical finance, are general and can be applied to Gibrat's law and many of its extensions in economics and finance.[1] According to our general characterization, any stationary distribution in a random growth setting is shaped entirely by two factors—the idiosyncratic volatilities and reversion rates (a measure of cross-sectional mean reversion) for different ranks in the distribution. An increase in idiosyncratic volatilities increases concentration, while an increase in reversion rates decreases concentration. We also present results that allow for the estimation of these two factors using panel data.

Our characterization of a stationary distribution in a general, nonparametric setting provides a framework in which we can understand the shaping forces for almost all power law distributions that emerge in random growth settings. After all, one implication of our results is that the distributional effect of any economic mechanism can be inferred by determining the effect of that mechanism on the idiosyncratic volatilities and reversion rates for different ranked processes. In addition to our general characterization, this is the first paper in economics to provide empirical methods to measure the econometric factors that shape power law distributions using panel data.

These empirical methods allow us to understand the causes, in an econometric sense, of

[1]There is a growing and extensive literature analyzing these rank-based methods. See, for example, Banner et al. (2005), Pal and Pitman (2008), Ichiba et al. (2011), and Shkolnikov (2011).

distributional changes that occur for power laws in economics and finance. This econometric analysis has many potential applications. For example, our methods can establish that increasing U.S. income and wealth inequality (Atkinson et al., 2011; Saez and Zucman, 2014) is the result of changes in either reversion rates or the magnitude of idiosyncratic shocks to household income and wealth. These econometric changes could then be linked to the evolution of policy, skill-biased technological change, or other changes in the economic environment. This is the approach of Fernholz and Koch (2016), who analyze the increasing concentration of U.S. bank assets using the econometric techniques presented in this paper. A similar analysis of increasing U.S. house price dispersion (Van Nieuwerburgh and Weill, 2010) should also yield new conclusions and useful insight.

In order to demonstrate the validity and accuracy of our empirical methods, we estimate reversion rates and idiosyncratic volatilities using monthly commodity prices data from 1980 - 2015 and compare the predicted distribution of relative commodity prices using our methods to the average distribution of relative commodity prices observed during this period. Although our methods apply most naturally to distributions such as wealth, firm size, and city size, they can also be applied to the distributions of relative asset prices. By testing our methods using normalized commodity prices data, we are able to examine the applicability of these methods to relative asset price distributions.

Commodity prices must be normalized so that they can be compared in an economically meaningful way, but as long as these appropriately-normalized prices satisfy the basic regularity conditions that our econometric theory relies on, then our methods should be applicable to the distribution of relative commodity prices. Furthermore, because the distribution of relative normalized commodity prices appears to be stationary during the 1980 - 2015 period, the rank-based reversion rates and idiosyncratic volatilities that we estimate should provide an accurate description of the observed relative commodity price distribution during this period. We confirm that this is in fact the case. One of the contributions of this paper, then, is to show that our empirical methods can validly be applied not only to standard size distributions but also to relative asset price distributions. This result highlights the potential for future applications of our econometric techniques using other data sets.

In addition to our characterization of a stationary distribution in a general random growth setting, we also show that a mean-reversion condition is necessary for the existence of such a stationary distribution. Specifically, a stationary distribution exists only if the growth

rates of higher-ranked processes are on average lower than the growth rates of lower-ranked processes. If we let the processes in our general random growth setting represent the total market capitalizations of different stocks, then this mean-reversion condition implies that bigger stocks must generate lower capital gains than smaller stocks. This is similar to the well-known size effect for stocks—the tendency for U.S. stocks with large total market capitalizations to generate lower average returns than U.S. stocks with small total market capitalizations (Banz, 1981).

In terms of normalized commodity prices, the mean-reversion condition that we describe as necessary for the existence of a stationary distribution provides a testable prediction of the existence of a generalized "size" effect for commodities. That is, our results predict that higher-ranked, higher-priced, "bigger" commodities should generate lower returns on average than lower-ranked, lower-priced, "smaller" commodities. Using the same monthly commodity prices data from 1980 - 2015, we confirm that this is in fact the case. We show that an equal-weighted portfolio that invests only in the most expensive (highest ranked) commodities each month generates a yearly return on average more than 6% below the return on an equal-weighted portfolio that invests only in the least expensive (lowest ranked) commodities each month, exactly as predicted by our econometric theory. Future research that analyzes the risk and liquidity properties of this generalized "size" effect for commodities and attempts to determine if this excess return is consistent with standard equilibrium asset pricing theories (Lucas, 1978; Fama and French, 1993) could yield interesting results.

The rest of this paper is organized as follows. Section 2 presents our nonparametric framework and derives the main result that characterizes general stationary power law distributions. Section 3 presents results that show how to estimate the two shaping factors of a power law distribution using panel data. Section 4 presents estimates of rank-based reversion rates and idiosyncratic volatilities using commodity prices data, and also shows the existence of a large generalized "size" effect as predicted by our econometric results. Section 5 concludes. Appendix A discusses the regularity assumptions needed for our main results, and Appendix B contains all proofs.

2 A Nonparametric Approach to Dynamic Power Law Distributions

For consistency, we shall refer to agents holding units throughout this section. However, it is important to note that in this general setup agents can represent households, firms, cities, countries, and other entities, with the corresponding units representing income, wealth, total employees, population, and other quantities. Furthermore, we can also interpret agents' holdings of units as the prices of different assets, as we shall do for commodity prices in Section 4 below.

Consider a population that consists of $N > 1$ agents. Time is continuous and denoted by $t \in [\,0, \infty)$, and uncertainty in this population is represented by a filtered probability space $(\Omega, \mathcal{F}, \mathcal{F}_t, P)$. Let $\mathbf{B}(t) = (B_1(t), \ldots, B_M(t))$, $t \in [0, \infty)$, be an M-dimensional Brownian motion defined on the probability space, with $M \geq N$. We assume that all stochastic processes are adapted to $\{\mathcal{F}_t; t \in [0, \infty)\}$, the augmented filtration generated by $\mathbf{B}$.[2]

2.1 Dynamics

The total units held by each agent $i = 1, \ldots, N$ is given by the process x_i. Each of these unit processes evolves according to the stochastic differential equation

$$d \log x_i(t) = \mu_i(t)\, dt + \sum_{s=1}^{M} \delta_{is}(t)\, dB_s(t), \tag{2.1}$$

where μ_i and δ_{is}, $s = 1, \ldots, M$, are measurable and adapted processes. The growth rates and volatilities, μ_i and δ_{is}, respectively, are general and practically unrestricted, having only to satisfy a few basic regularity conditions that are discussed in Appendix A. These conditions imply that the unit processes for the agents are continuous semimartingales, which represent a broad class of stochastic processes (for a detailed discussion, see Karatzas and Shreve, 1991).[3]

Indeed, the martingale representation theorem (Nielsen, 1999) implies that any plausible

[2] In order to simplify the exposition, we shall omit many of the less important regularity conditions and technical details involved with continuous-time stochastic processes.

[3] This basic setup shares much in common with the continuous-time finance literature (see, for example, Karatzas and Shreve, 1998; Duffie, 2001). Continuous semimartingales are more general than Itô processes, which are common in the continuous-time finance literature (Nielsen, 1999).

continuous process for agents' unit holdings can be written in the nonparametric form of equation (2.1). Furthermore, this section's results can also apply to processes that are subject to sporadic, discontinuous jumps.[4] As a consequence, all previous analyses based on Gibrat's law or specific extensions to Gibrat's law (Gabaix, 1999, 2009) are special cases of our general framework in this paper.

It is useful to describe the dynamics of the total units held by all agents, which we denote by $x(t) = x_1(t) + \cdots + x_N(t)$. In order to do so, we first characterize the covariance of unit holdings across different agents over time. For all $i, j = 1, \ldots, N$, let the covariance process ρ_{ij} be given by

$$\rho_{ij}(t) = \sum_{s=1}^{M} \delta_{is}(t)\delta_{js}(t). \tag{2.2}$$

Applying Itô's Lemma to equation (2.1), we are now able to describe the dynamics of the total units process x.

Lemma 2.1. *The dynamics of the process for total units held by all agents x are given by*

$$d\log x(t) = \mu(t)\,dt + \sum_{i=1}^{N}\sum_{s=1}^{M} \theta_i(t)\delta_{is}(t)\,dB_s(t), \quad \text{a.s.}, \tag{2.3}$$

where

$$\theta_i(t) = \frac{x_i(t)}{x(t)}, \tag{2.4}$$

for $i = 1, \ldots, N$, and

$$\mu(t) = \sum_{i=1}^{N} \theta_i(t)\mu_i(t) + \frac{1}{2}\left(\sum_{i=1}^{N} \theta_i(t)\rho_{ii}(t) - \sum_{i,j=1}^{N} \theta_i(t)\theta_j(t)\rho_{ij}(t)\right). \tag{2.5}$$

2.2 Rank-Based Dynamics

In order to characterize the stationary distribution of units in this setup, it is necessary to consider the dynamics of agents' unit holdings by rank. One of the key insights of our approach and of this paper more generally is that rank-based unit dynamics are the essential determinants of the distribution of units. As we demonstrate below, there is a simple, direct, and robust relationship between rank-based unit growth rates and the distribution of units.

[4]This is an open area for research, but such extensions are examined by Shkolnikov (2011) and Fernholz (2016a).

This relationship is a purely statistical result and hence can be applied to essentially any economic environment, no matter how complex.

The first step in achieving this characterization is to introduce notation for agent rank based on unit holdings. For $k = 1, \ldots, N$, let

$$x_{(k)}(t) = \max_{1 \leq i_1 < \cdots < i_k \leq N} \min\left(x_{i_1}(t), \ldots, x_{i_k}(t)\right), \tag{2.6}$$

so that $x_{(k)}(t)$ represents the units held by the agent with the k-th most units among all the agents in the population at time t. For brevity, we shall refer to this agent as the k-th largest agent throughout this paper. One consequence of this definition is that

$$\max(x_1(t), \ldots, x_N(t)) = x_{(1)}(t) \geq x_{(2)}(t) \geq \cdots \geq x_{(N)}(t) = \min(x_1, \ldots, x_N(t)). \tag{2.7}$$

Next, let $\theta_{(k)}(t)$ be the share of total units held by the k-th largest agent at time t, so that

$$\theta_{(k)}(t) = \frac{x_{(k)}(t)}{x(t)}, \tag{2.8}$$

for $k = 1, \ldots, N$.

The next step is to describe the dynamics of the agent rank unit processes $x_{(k)}$ and rank unit share processes $\theta_{(k)}$, $k = 1, \ldots, N$. Unfortunately, this task is complicated by the fact that the max and min functions from equation (2.6) are not differentiable, and hence we cannot simply apply Itô's Lemma in this case. Instead, we introduce the notion of a local time to solve this problem. For any continuous process z, the *local time* at 0 for z is the process Λ_z defined by

$$\Lambda_z(t) = \frac{1}{2}\left(|z(t)| - |z(0)| - \int_0^t \operatorname{sgn}(z(s))\, dz(s)\right). \tag{2.9}$$

As detailed by Karatzas and Shreve (1991), the local time for z measures the amount of time the process z spends near zero.[5] To be able to link agent rank to agent index, let p_t be the random permutation of $\{1, \ldots, N\}$ such that for $1 \leq i, k \leq N$,

$$p_t(k) = i \quad \text{if} \quad x_{(k)}(t) = x_i(t). \tag{2.10}$$

[5]For more discussion of local times, and especially their connection to rank processes, see Fernholz (2002).

This definition implies that $p_t(k) = i$ whenever agent i is the k-th largest agent in the population at time t, with ties broken in some consistent manner.[6]

Lemma 2.2. *For all $k = 1, \ldots, N$, the dynamics of the agent rank unit processes $x_{(k)}$ and rank unit share processes $\theta_{(k)}$ are given by*

$$d\log x_{(k)}(t) = d\log x_{p_t(k)}(t) + \frac{1}{2}d\Lambda_{\log x_{(k)} - \log x_{(k+1)}}(t) - \frac{1}{2}d\Lambda_{\log x_{(k-1)} - \log x_{(k)}}(t), \tag{2.11}$$

a.s, and

$$d\log \theta_{(k)}(t) = d\log \theta_{p_t(k)}(t) + \frac{1}{2}d\Lambda_{\log \theta_{(k)} - \log \theta_{(k+1)}}(t) - \frac{1}{2}d\Lambda_{\log \theta_{(k-1)} - \log \theta_{(k)}}(t), \tag{2.12}$$

a.s., with the convention that $\Lambda_{\log x_{(0)} - \log x_{(1)}}(t) = \Lambda_{\log x_{(N)} - \log x_{(N+1)}}(t) = 0$.

According to equation (2.11) from the lemma, the dynamics of units for the k-th largest agent in the population are the same as those for the agent that is the k-th largest at time t (agent $i = p_t(k)$), plus two local time processes that capture changes in agent rank (one agent overtakes another in unit holdings) over time.[7] Equation (2.12) describes the similar dynamics of the rank unit share processes $\theta_{(k)}$.

Using equations (2.1) and (2.3) and the definition of $\theta_i(t)$, we have that for all $i = 1, \ldots, N$,

$$\begin{aligned} d\log \theta_i(t) &= d\log x_i(t) - d\log x(t) \\ &= \mu_i(t)\,dt + \sum_{s=1}^{M} \delta_{is}(t)\,dB_s(t) - \mu(t)\,dt - \sum_{i=1}^{N}\sum_{s=1}^{M} \theta_i(t)\delta_{is}(t)\,dB_s(t). \end{aligned} \tag{2.13}$$

If we apply Lemma 2.2 to equation (2.13), then it follows that

$$\begin{aligned} d\log \theta_{(k)}(t) = \big(\mu_{p_t(k)}(t) - \mu(t)\big)\,dt + \sum_{s=1}^{M} \delta_{p_t(k)s}(t)\,dB_s(t) - \sum_{i=1}^{N}\sum_{s=1}^{M} \theta_i(t)\delta_{is}(t)\,dB_s(t) \\ + \frac{1}{2}d\Lambda_{\log \theta_{(k)} - \log \theta_{(k+1)}}(t) - \frac{1}{2}d\Lambda_{\log \theta_{(k-1)} - \log \theta_{(k)}}(t), \end{aligned} \tag{2.14}$$

a.s, for all $k = 1, \ldots, N$. Equation (2.14), in turn, implies that the process $\log \theta_{(k)} - \log \theta_{(k+1)}$

[6]For example, if $x_i(t) = x_j(t)$ and $i > j$, then we can set $p_t(k) = i$ and $p_t(k+1) = j$.

[7]For brevity, we write $dz_{p_t(k)}(t)$ to refer to the process $\sum_{i=1}^{N} 1_{\{i=p_t(k)\}} dz_i(t)$ throughout this paper.

satisfies, a.s., for all $k = 1, \ldots, N-1$,

$$\begin{aligned} d\left(\log\theta_{(k)}(t) - \log\theta_{(k+1)}(t)\right) = &\left(\mu_{p_t(k)}(t) - \mu_{p_t(k+1)}(t)\right)\, dt + d\Lambda_{\log\theta_{(k)}-\log\theta_{(k+1)}}(t) \\ &- \frac{1}{2} d\Lambda_{\log\theta_{(k-1)}-\log\theta_{(k)}}(t) - \frac{1}{2} d\Lambda_{\log\theta_{(k+1)}-\log\theta_{(k+2)}}(t) \\ &+ \sum_{s=1}^{M} \left(\delta_{p_t(k)s}(t) - \delta_{p_t(k+1)s}(t)\right)\, dB_s(t). \end{aligned} \tag{2.15}$$

The processes for relative unit holdings of adjacent agents in the distribution of units as given by equation (2.15) are key to describing the distribution of units in this setup.

2.3 Stationary Distribution

The results presented above allow us to analytically characterize the stationary distribution of units in this setup. Let α_k equal the time-averaged limit of the expected growth rate of units for the k-th largest agent relative to the expected growth rate of units for the entire population of agents, so that

$$\alpha_k = \lim_{T\to\infty} \frac{1}{T} \int_0^T \left(\mu_{p_t(k)}(t) - \mu(t)\right)\, dt, \tag{2.16}$$

for $k = 1, \ldots, N$. The relative growth rates α_k are a rough measure of the rate at which agents' unit holdings revert to the mean. We shall refer to the $-\alpha_k$ as reversion rates, since lower values of α_k (and hence higher values of $-\alpha_k$) imply faster cross-sectional mean reversion.

In a similar manner, we wish to define the time-averaged limit of the volatility of the process $\log\theta_{(k)} - \log\theta_{(k+1)}$, which measures the relative unit holdings of adjacent agents in the distribution of units. For all $k = 1, \ldots, N-1$, let σ_k be given by

$$\sigma_k^2 = \lim_{T\to\infty} \frac{1}{T} \int_0^T \sum_{s=1}^{M} \left(\delta_{p_t(k)s}(t) - \delta_{p_t(k+1)s}(t)\right)^2\, dt. \tag{2.17}$$

The relative growth rates α_k together with the volatilities σ_k entirely determine the shape of the stationary distribution of units in this population, as we shall demonstrate below.

We shall refer to the volatility parameters σ_k, which measure the standard deviations of the processes $\log\theta_{(k)} - \log\theta_{(k+1)}$, as idiosyncratic volatilities. An idiosyncratic shock to the

unit holdings of either the k-th or $(k+1)$-th ranked agent alters the value of $\log\theta_{(k)}-\log\theta_{(k+1)}$ and hence will be measured by σ_k. In addition, however, a shock that affects the unit holdings of multiple agents that do not occupy adjacent ranks in the distribution will also alter this value. Indeed, any shock that affects $\log\theta_{(k)}$ and $\log\theta_{(k+1)}$ differently, must necessarily alter the value of $\log\theta_{(k)}-\log\theta_{(k+1)}$ and hence will be measured by σ_k. In this sense, the volatility parameters σ_k are slightly more general than pure idiosyncratic volatilities that capture only shocks that affect one single agent at a time.

Finally, for all $k=1,\ldots,N$, let

$$\kappa_k=\lim_{T\to\infty}\frac{1}{T}\Lambda_{\log\theta_{(k)}-\log\theta_{(k+1)}}(T). \tag{2.18}$$

Let $\kappa_0=0$, as well. Throughout this paper, we assume that the limits in equations (2.16)-(2.18) do in fact exist. In Appendix B, we show that the parameters α_k and κ_k are related by $\alpha_k-\alpha_{k+1}=\frac{1}{2}\kappa_{k-1}-\kappa_k+\frac{1}{2}\kappa_{k+1}$, for all $k=1,\ldots,N-1$.

The *stable version* of the process $\log\theta_{(k)}-\log\theta_{(k+1)}$ is the process $\log\hat{\theta}_{(k)}-\log\hat{\theta}_{(k+1)}$ defined by

$$d\left(\log\hat{\theta}_{(k)}(t)-\log\hat{\theta}_{(k+1)}(t)\right)=-\kappa_k\,dt+d\Lambda_{\log\hat{\theta}_{(k)}-\log\hat{\theta}_{(k+1)}}(t)+\sigma_k\,dB(t), \tag{2.19}$$

for all $k=1,\ldots,N-1$.[8] The stable version of $\log\theta_{(k)}-\log\theta_{(k+1)}$ replaces all of the processes from the right-hand side of equation (2.15) with their time-averaged limits, with the exception of the local time process $\Lambda_{\log\theta_{(k)}-\log\theta_{(k+1)}}$. By considering the stable version of these relative unit holdings processes, we are able to obtain a simple characterization of the distribution of units.

Theorem 2.3. *There is a stationary distribution for the stable version of unit holdings by agents in this population if and only if $\alpha_1+\cdots+\alpha_k<0$, for $k=1,\ldots,N-1$. Furthermore, if there is a stationary distribution of units, then for $k=1,\ldots,N-1$, this distribution satisfies*

$$E\left[\log\hat{\theta}_{(k)}(t)-\log\hat{\theta}_{(k+1)}(t)\right]=\frac{\sigma_k^2}{-4(\alpha_1+\cdots+\alpha_k)},\quad\text{a.s.} \tag{2.20}$$

Theorem 2.3 provides an analytic rank-by-rank characterization of the entire distribution of units. This is achieved despite minimal assumptions on the processes that describe the

[8]For each $k=1,\ldots,N-1$, equation (2.19) implicitly defines another Brownian motion $B(t)$, $t\in[0,\infty)$. These Brownian motions can covary in any way across different k.

dynamics of agents' unit holdings over time. As long as the relative growth rates, volatilities, and local times that we take limits of in equations (2.16)-(2.18) do not change drastically and frequently over time, then the distribution of the stable versions of $\theta_{(k)}$ from Theorem 2.3 will accurately reflect the distribution of the true versions of these rank unit share processes.[9] For this reason, we shall assume that equation (2.20) approximately describes the true versions of $\theta_{(k)}$ throughout much of this paper.

The theorem yields two important insights. First, it shows that an understanding of rank-based unit holdings dynamics is sufficient to describe the entire distribution of units. It is not necessary to directly model and estimate agents' unit holdings dynamics by name, denoted by index i, as is common in the literatures on income and wealth inequality (Guvenen, 2009; Benhabib et al., 2011; Altonji et al., 2013). Second, the theorem shows that the only two factors that affect the distribution of units are the rank-based reversion rates, $-\alpha_k$, and the rank-based volatilities, σ_k.

The characterization in equation (2.20) is flexible enough to replicate any empirical distribution. Indeed, regardless of whether the true distribution of units is Pareto, log-normal, double Pareto log-normal, or something else, Theorem 2.3 implies that this distribution appears asymptotically for certain values of the reversion rates and volatilities.

According to Theorem 2.3, stationarity of the distribution of agents' unit holdings requires that the reversion rates $-\alpha_k$ must sum to positive quantities, for all $k = 1, \ldots, N-1$. Stability, then, requires a mean reversion condition in the sense that the growth rate of units for the agents with the most units in the population must be strictly below the growth rate of units for agents with smaller unit holdings. The unstable case in which this mean reversion condition does not hold is examined in detail by Fernholz and Fernholz (2014) and Fernholz (2016b). As we shall demonstrate in Section 4, this condition has significant implications for the dynamics of different ranked commodity prices.

If we impose more restrictions on the stable versions of the relative unit holdings processes $\log \hat{\theta}_{(k)} - \log \hat{\theta}_{(k+1)}$, then it is possible to link the reversion rates $-\alpha_k$ and volatilities σ_k to mobility. In particular, if we assume that agents face only aggregate and idiosyncratic shocks to their unit holdings, then it is possible to show that mobility is increasing in cross-sectional mean reversion $-\alpha_k$ and decreasing in unit concentration, as measured by the expected value

[9] Fernholz (2002) and Fernholz and Koch (2016) demonstrate the accuracy of Theorem 2.3 in matching, respectively, the distribution of total market capitalizations of U.S. stocks and the distribution of assets of U.S. financial intermediaries.

of $\log \hat{\theta}_{(k)} - \log \hat{\theta}_{(k+1)}$. Mobility in this context is measured as the expected time for $x_{(k+1)}$ to overtake the higher ranked $x_{(k)}$. A proof of this result and some extensions can be found in Fernholz (2016a).

2.4 Gibrat's Law, Zipf's Law, and Pareto Distributions

It is useful to see how our rank-based, nonparametric approach nests many common examples of random growth processes from other literatures as special cases. We shall focus on the influential example of Gibrat's law, and also describe the conditions that are necessary for Gibrat's law to give rise to Zipf's law.

According to Gabaix (2009), the strongest form of Gibrat's law for unit holdings imposes growth rates and volatilities that do not vary across the distribution of unit holdings. In terms of the reversion rates $-\alpha_k$ (which measure relative unit growth rates for different ranked agents) and idiosyncratic volatilities σ_k, this requirement is equivalent to there existing some common $\alpha < 0$ and $\sigma > 0$ such that

$$\alpha = \alpha_1 = \cdots = \alpha_{N-1}, \tag{2.21}$$

and

$$\sigma = \sigma_1 = \cdots = \sigma_{N-1}. \tag{2.22}$$

In terms of equation (2.20) from Theorem 2.3, then, Gibrat's law yields unit shares that satisfy

$$E\left[\log \hat{\theta}_{(k)}(t) - \log \hat{\theta}_{(k+1)}(t)\right] = \frac{\sigma_k^2}{-4(\alpha_1 + \cdots + \alpha_k)} = \frac{\sigma^2}{-4k\alpha} \quad \text{a.s.}, \tag{2.23}$$

for all $k = 1, \ldots, N-1$.

The distribution of agents' unit holdings follows a Pareto distribution if a plot of unit shares as a function of rank, using log scales for both axes, appears as a straight line.[10] Furthermore, if the slope of such a straight line plot is -1, then agents' unit shares obey Zipf's law (Gabaix, 1999). According to equation (2.23), for all $k = 1, \ldots, N-1$, the slope

[10] See the discussions in Newman (2006) and Gabaix (2009).

of such a log-log plot in the case of Gibrat's law is given by

$$\frac{E\left[\log\hat{\theta}_{(k)}(t)-\log\hat{\theta}_{(k+1)}(t)\right]}{\log k-\log k+1}\approx -kE\left[\log\hat{\theta}_{(k)}(t)-\log\hat{\theta}_{(k+1)}(t)\right]=\frac{-k\sigma^2}{-4k\alpha}=\frac{\sigma^2}{4\alpha}. \quad (2.24)$$

Equation (2.24) shows that Gibrat's law yields a Pareto distribution in which the log-log plot of unit shares versus rank has slope $\sigma^2/4\alpha<0$, which is equivalent to the Pareto distribution having parameter $-\sigma^2/4\alpha>0$. Furthermore, we see that agents' unit shares obey Zipf's law only if $\sigma^2=-4\alpha$, in which case the log-log plot has slope -1.

Theorem 2.3 thus demonstrates that Gibrat's law and Zipf's law are special cases of general power law distributions in which growth rates and volatilities potentially vary across different ranks in the distribution of unit holdings. Indeed, equation (2.20) implies that any power law exponent can obtain in any part of the distribution curve. This flexibility is a novel feature of our empirical methodology and is necessary to accurately match many empirical distributions. For example, Fernholz and Koch (2016) find that asset growth rates and volatilities vary substantially across different size-ranked U.S. financial intermediaries. Similarly, Fernholz (2002) finds that growth rates and volatilities of total market capitalization vary substantially across different size-ranked U.S. stocks, while Neumark et al. (2011) find that employment growth rates vary across different size-ranked U.S. firms. In Section 4, we confirm this general pattern and show that the growth rates of commodity prices also differ across ranks in a statistically significant and economically meaningful way.

3 Estimation

In order to estimate the reversion rates and volatilities from equation (2.20) from Theorem 2.3, we use discrete-time approximations of the continuous processes that yield the theorem. For the estimation of the volatility parameters σ_k^2, we use the discrete-time approximation of equation (2.17) above. In particular, these estimates are given by

$$\sigma_k^2=\frac{1}{T}\sum_{t=1}^{T}\left[\left(\log\theta_{p_t(k)}(t+1)-\log\theta_{p_t(k+1)}(t+1)\right)-\left(\log\theta_{p_t(k)}(t)-\log\theta_{p_t(k+1)}(t)\right)\right]^2, \quad (3.1)$$

for all $k = 1, \ldots, N-1$. Note that T is the total number of periods covered in the data.

The estimation of the rank-based relative growth rates α_k is more difficult. In order to estimate these parameters, we first estimate the local time parameters κ_k and then exploit the relationship that exists between these local times and the rank-based relative growth rates.

Lemma 3.1. *The relative growth rate parameters α_k and the local time parameters κ_k satisfy*

$$\alpha_k = \frac{1}{2}\kappa_{k-1} - \frac{1}{2}\kappa_k, \tag{3.2}$$

for all $k = 1, \ldots, N-1$, and $\alpha_N = -(\alpha_1 + \cdots + \alpha_{N-1})$.

Lemma 3.2. *The ranked agent unit share processes $\theta_{(k)}$ satisfy the stochastic differential equation*

$$\begin{aligned} d\log\left(\theta_{p_t(1)}(t) + \cdots + \theta_{p_t(k)}(t)\right) &= d\log\left(\theta_{(1)}(t) + \cdots + \theta_{(k)}(t)\right) \\ &\quad - \frac{\theta_{(k)}(t)}{2(\theta_{(1)}(t) + \cdots + \theta_{(k)}(t))} d\Lambda_{\log\theta_{(k)} - \log\theta_{(k+1)}}(t), \quad \text{a.s.}, \end{aligned} \tag{3.3}$$

for all $k = 1, \ldots, N$.

These lemmas together allow us to generate estimates of the rank-based relative growth rates α_k. In order to accomplish this, we first estimate the local time processes $\Lambda_{\log\hat{\theta}_{(k)} - \log\hat{\theta}_{(k+1)}}$ using the discrete-time approximation of equation (3.3). This discrete-time approximation implies that for all $k = 1, \ldots, N$,

$$\begin{aligned} &\log\left(\theta_{p_t(1)}(t+1) + \cdots + \theta_{p_t(k)}(t+1)\right) - \log\left(\theta_{p_t(1)}(t) + \cdots + \theta_{p_t(k)}(t)\right) = \\ &\qquad \log\left(\theta_{p_{t+1}(1)}(t+1) + \cdots + \theta_{p_{t+1}(k)}(t+1)\right) - \log\left(\theta_{p_t(1)}(t) + \cdots + \theta_{p_t(k)}(t)\right) \\ &\qquad\quad - \frac{\theta_{p_t(k)}(t)}{2\left(\theta_{p_t(1)}(t) + \cdots + \theta_{p_t(k)}(t)\right)} \left(\Lambda_{\log\theta_{(k)} - \log\theta_{(k+1)}}(t+1) - \Lambda_{\log\theta_{(k)} - \log\theta_{(k+1)}}(t)\right), \end{aligned} \tag{3.4}$$

which, after simplification and rearrangement, yields

$$\Lambda_{\log\theta_{(k)}-\log\theta_{(k+1)}}(t+1)-\Lambda_{\log\theta_{(k)}-\log\theta_{(k+1)}}(t)=\Big[\log\left(\theta_{p_{t+1}(1)}(t+1)+\cdots+\theta_{p_{t+1}(k)}(t+1)\right) - \log\left(\theta_{p_t(1)}(t+1)+\cdots+\theta_{p_t(k)}(t+1)\right)\Big]\frac{2\left(\theta_{p_t(1)}(t)+\cdots+\theta_{p_t(k)}(t)\right)}{\theta_{p_t(k)}(t)}. \tag{3.5}$$

As with our estimates of the volatility parameters σ_k^2, we estimate the values of the local times in equation (3.5) for $t=1,\ldots,T$, where T is the total number of periods covered in the data. We also set $\Lambda_{\log\theta_{(k)}-\log\theta_{(k+1)}}(0)=0$, for all $k=1,\ldots,N$.

After estimating the local times in equation (3.5), we then use equation (2.18) to generate estimates of κ_k according to

$$\kappa_k=\frac{1}{T}\Lambda_{\log\theta_{(k)}-\log\theta_{(k+1)}}(T), \tag{3.6}$$

for all $k=1,\ldots,N$. Finally, we can use the relationship between the parameters α_k and κ_k established by Lemma 3.1. This is accomplished via equation (3.2), which yields estimates of each α_k using our estimates of the parameters κ_k from equation (3.6).

While the methods described in this section explain how to generate point estimates of the reversion rates $-\alpha_k$ and idiosyncratic volatilities σ_k, it is important to also understand how much variation there is in these estimates. It is not possible to generate confidence intervals using classical techniques in this setting because the empirical distribution of the parameters α_k and σ_k is unknown. However, it is possible to use bootstrap resampling to generate confidence intervals for these estimated factors.

Equations (3.1) and (3.5) show that the reversion rates $-\alpha_k$ and idiosyncratic volatilities σ_k are measured as changes from one period, t, to the next, $t+1$. As a consequence, the bootstrap resamples we construct consist of $T-1$ pairs of observations of agents' unit holdings from adjacent time periods (periods t and $t+1$). Such resamples, of course, are equivalent to the full sample which has observations over T periods and hence consists of $T-1$ pairs of observations from adjacent periods. The confidence intervals are then generated by determining the range of values that obtain for the parameters α_k and σ_k over all of the bootstrap resamples. In Section 4, we apply our techniques to the distribution of relative commodity prices and generate confidence intervals for our estimates of the reversion rates $-\alpha_k$ and idiosyncratic volatilities σ_k following this procedure.

4 Application: The Distribution of Commodity Prices

We wish to confirm the validity and accuracy of the empirical methods we presented in Section 2. We do this using a publicly available data set on the global monthly spot prices of 22 common commodities for 1980 - 2015 obtained from the Federal Reserve Bank of St. Louis (FRED).[11]

In order to accomplish this, we shall use the results and procedure described in Section 3 to estimate rank-based reversion rates $-\alpha_k$ and idiosyncratic volatilities σ_k for the distribution of relative commodity prices over our sample period 1980 - 2015. In this section, then, we shall interpret agents' holdings of units $x_i(t)$ from equation (2.1) as the prices of different commodities.[12] Because commodities are sold in different units and hence their prices cannot be compared in an economically meaningful way, it is important to normalize these prices by equalizing them in the initial period.

The results of Sections 2 and 3 apply to the distribution of the parameters $\theta_{(k)}$, $k = 1, \ldots, N$, which in those sections represented the shares of total units held by different ranked agents. If we interpret the x_i as commodity prices, then the parameters $\theta_{(k)}$ represent commodity "price shares," a quantity that is well defined but difficult to interpret economically. It is easy to show, however, that the distribution of these commodity "price shares" $\theta_{(k)}$ is the same as the distribution of commodity prices relative to the average of all commodity prices. This latter quantity has a clear economic interpretation. In this section, we estimate reversion rates and idiosyncratic volatilities that describe the stationary distribution of relative normalized commodity prices according to equation (2.20). As Figure 1 demonstrates, the distribution of these relative normalized prices appears to be roughly stationary over time. Consistent with this observation, we confirm below that the methods presented in Sections 2 and 3 do in fact accurately describe this stationary distribution.

If we let $\bar{x}(t)$ equal the average price of all N commodities at time t, then for all $i = 1, \ldots, N$, the relative price of commodity i at time t is defined as

$$\tilde{x}_i(t) = \frac{x_i(t)}{\bar{x}(t)} = \frac{x_i(t)}{\frac{x_1(t)+\cdots+x_N(t)}{N}} = \frac{Nx_i(t)}{x_1(t)+\cdots+x_N(t)}. \tag{4.1}$$

[11] These commodities are aluminum, bananas, barley, beef, Brent crude oil, cocoa, copper, corn, cotton, iron, lamb, lead, nickel, orange, poultry, rubber, soybeans, sugar, tin, wheat, wool (fine), and zinc.

[12] As long as commodity prices satisfy the basic regularity conditions of Appendix A and the distribution of relative commodity prices is stationary, then the econometric results of Sections 2 and 3 can be applied.

The relative price $\tilde{x}_i(t)$ is equal to the price of commodity i at time t relative to the average price of all N commodities at time t. If we let $\tilde{x}_{(k)}(t)$ denote the relative price of the k-th ranked commodity at time t, then equations (2.4) and (2.8) imply that, for all $i, k = 1, \ldots, N$,

$$\tilde{x}_i(t) = N\theta_i(t) \qquad \text{and} \qquad \tilde{x}_{(k)}(t) = N\theta_{(k)}(t). \tag{4.2}$$

Note that the k-th ranked commodity at time t refers to the commodity with the k-th highest price at time t. It follows from equation (4.2) that for all $k = 1, \ldots, N-1$ and all t,

$$\log \tilde{x}_{(k)}(t) - \log \tilde{x}_{(k+1)}(t) = \log \theta_{(k)}(t) - \log \theta_{(k+1)}(t), \tag{4.3}$$

and hence equation (2.20) from Theorem 2.3 describes both the distribution of commodity "price shares," $\theta_{(k)}$, and relative commodity prices, $\tilde{x}_{(k)}$. In other words, all of our previous results apply to the distribution of relative commodity prices as well.

4.1 Prediction and Data

The econometric results of Section 2 suggest that any stationary size distribution can be accurately characterized by the reversion rates $-\alpha_k$ and idiosyncratic volatilities σ_k according to equation (2.20). Fernholz (2002) and Fernholz and Koch (2016) show, respectively, that this is in fact true for the size distributions of total market capitalizations of U.S. stocks and total assets of U.S. financial intermediaries. One of this paper's contributions is to further demonstrate the validity of our econometric techniques using a new data set.

The first step is to estimate the reversion rates $-\alpha_k$ for each rank $k = 1, \ldots, N$. As described in Section 2, these reversion rates measure the growth rates of different ranked commodity prices relative to the growth rate of all commodity prices together. In Figure 2, we plot annualized values of minus the reversion rates α_k for each rank in the distribution of relative normalized commodity prices together with 95% confidence intervals based on the results of 10,000 bootstrap resample estimates.

These parameters are estimated using the procedure described in Section 3. In particular, these estimated reversion rates are generated by first estimating the local time parameters κ_k according to equations (3.5) and (3.6), and then generating estimates of the parameters α_k according to equation (3.2) from Lemma 3.1. Figure 3 plots the evolution of the local time processes $\Lambda_{\log \tilde{x}_{(k)} - \log \tilde{x}_{(k+1)}}$, $k = 1, \ldots, N-1$, which we use to construct our estimates

of the reversion rates $-\alpha_k$.

The confidence intervals in Figure 2 are generated using this same procedure, only with bootstrap resamples instead of the original full sample. These confidence intervals show that the deviations from Gibrat's law for commodity prices during the 1980 - 2015 period are highly statistically significant. This observation confirms the usefulness of our rank-based methods, since these methods allow for growth rates that vary across the distribution of relative commodity prices in the realistic manner shown in Figure 2.

The next step is to estimate the idiosyncratic volatilities σ_k, which is accomplished using the discrete-time approximation given by equation (3.1). Figure 4 plots annualized estimates for these parameter values for each rank in the distribution of relative normalized commodity prices together with 95% confidence intervals based on the results of 10,000 bootstrap resample estimates. The estimates and confidence intervals for the parameters α_k and σ_k in Figures 2 and 4 are smoothed across different ranks using a Gaussian kernel smoother. Following Fernholz and Koch (2016), we smooth these parameters between 1 and 100 times and then choose the number of smoothings within this range that minimizes the squared deviation between the predicted relative commodity prices according to equation (2.20) and the average observed relative commodity prices for the period 1981-2015.[13]

How well do the reversion rates $-\alpha_k$ and idiosyncratic volatilities σ_k reported in Figures 2 and 4 replicate the true distribution of relative commodity prices? Figure 5 shows that these estimated parameters generate predicted relative commodity prices according to equation (2.20) that do in fact match the average relative commodity prices observed during the 1980 - 2015 sample period. The squared deviation between predicted and observed average relative commodity prices over this sample period is 0.143. Thus, we further confirm the validity of our econometric methods using commodity prices data.

4.2 A "Size" Effect for Commodities

One of the implications of Theorem 2.3 is that there is a stationary distribution of relative commodity prices if and only if $\alpha_1 + \cdots + \alpha_k < 0$. In other words, the growth rates of the prices of the higher-priced, higher-ranked commodities must on average be lower than the

[13]The commodity prices are normalized to all equal each other at the start of our sample period in 1980. Since it takes a number of months for these initially equal relative prices to converge to a stationary distribution, we remove the first year of data when generating observed average relative prices to compare to predicted relative prices for the purposes of smoothing the parameters α_k and σ_k.

growth rates of the prices of the lower-priced, lower-ranked commodities, otherwise there is no stationary distribution of relative commodity prices. This necessary condition is essentially a mean-reversion condition.

Suppose that we interpret the processes x_i in equation (2.1) from Section 2 as the total market capitalizations of stocks. In this case, the dynamics of the processes x_i correspond to capital gains, and hence the mean-reversion condition from Theorem 2.3 implies that bigger stocks must generate smaller capital gains than smaller stocks. In other words, the mean-reversion condition from Theorem 2.3 offers a structural, econometric explanation for the well-known size effect for stocks—the tendency for U.S. stocks with large total market capitalizations to generate lower average returns than U.S. stocks with small total market capitalizations (Banz, 1981; Fama and French, 1993). Indeed, this condition implies that a long-run size effect for capital gains is a necessary consequence of a stationary and realistic distribution of total stock market capitalizations.

This surprising implication of Theorem 2.3 offers a testable prediction for our commodity prices data—there should be a generalized "size" effect for commodities in which higher-ranked, higher-priced, "bigger," commodities generate lower returns on average than lower-ranked, lower-priced, "smaller," commodities. In Figure 6, we plot the log values over time of a portfolio that invests equal quantities in the eleven most expensive commodities in each month and a portfolio that invests equal quantities in the eleven cheapest commodities in each month. More precisely, in each month t, the expensive commodities portfolio is rebalanced to invest an equal quantity of the portfolio value in month t in each of the eleven most expensive (highest ranked) commodities in month t. Conversely, in each month t, the cheap commodities portfolio is rebalanced to invest an equal quantity of the portfolio value in month t in each of the eleven cheapest (least expensive, lowest ranked) commodities in month t. Figure 6 demonstrates a clear and large generalized size effect for commodities, just as predicted by our econometric results in Section 2.

Figure 7 plots the log of the value of the cheap commodities portfolio relative to the value of the expensive commodities portfolio. This figure confirms the generalized size effect for commodities as in Figure 6. In terms of standard percentage returns, the cheap commodities portfolio generates an average yearly (monthly) return of 8.62% (0.59%), while the expensive commodities portfolio generates an average yearly (monthly) return of 2.25% (0.12%).

Figure 7 also shows that the excess return of the cheap commodities portfolio relative

to the expensive commodities portfolio does not appear to be highly positively correlated with either U.S. equity returns or the U.S. business cycle. This is surprising, since such positive correlations would be predicted from standard asset pricing theories (Lucas, 1978; Cochrane, 2005). Nonetheless, the generalized size effect for commodities predicted by the mean-reversion condition of Theorem 2.3 and confirmed in Figures 6 and 7 is not necessarily inconsistent with standard equilibrium asset pricing theories.

The mean-reversion condition of Theorem 2.3 implies that a generalized size effect for commodities is a necessary consequence of a realistic stationary distribution. It does not, however, imply anything about the properties of the excess returns from such a size effect. For example, the size effect shown in Figure 6 could be a reflection of greater risk for the portfolio of cheap commodities relative to the portfolio of expensive commodities, one possible explanation for the size effect among stocks (Fama and French, 1993). It could also be a reflection of lower liquidity for the portfolio of cheap commodities relative to the portfolio of expensive commodities, another possible explanation for the size effect among stocks (Acharya and Pedersen, 2005). The mean-reversion condition of Theorem 2.3 implies only that a generalized size effect for commodities is to be expected, regardless of whether or not such a size effect is a reflection of higher risk or lower liquidity.

It is beyond the scope of this paper to examine in detail the risk and liquidity properties of the two portfolio returns shown in Figure 6, but such an analysis may yield interesting insight about the asset-pricing implications of our econometric results. Furthermore, although we have confirmed the existence of a new generalized size effect for commodities as predicted by Theorem 2.3, the generality of our nonparametric, rank-based econometric framework in Section 2 suggests that generalized size effects should exist for other size and relative price distributions as well. As long as these other distributions are roughly stationary, then our theory predicts the existence of generalized size effects. Future research that attempts to uncover such new generalized size effects is likely to yield interesting conclusions.

5 Conclusion

This paper presents rank-based, nonparametric methods that allow for the characterization of general power law distributions in random growth settings. We show that any stationary distribution in a random growth setting is shaped entirely by two factors—the idiosyncratic volatilities and reversion rates (a measure of cross-sectional mean reversion) for different ranks in the distribution. An increase in idiosyncratic volatilities increases concentration, while an increase in reversion rates decreases concentration. We also provide methods for estimating these two shaping factors using panel data.

Using data on a set of 22 global commodity prices from 1980 - 2015, we show that our rank-based, nonparametric methods accurately describe the distribution of relative normalized commodity prices. According to our econometric results, a necessary condition for the existence of a stationary distribution is that higher ranked (more expensive) commodity prices must grow more slowly than lower ranked (less expensive) commodity prices. In other words, our results predict a generalized "size" effect for commodities in which lower-priced commodities generate higher returns than higher-priced commodities. We confirm this prediction and show that a portfolio of lower-priced commodities has substantially higher returns than a portfolio of higher-priced commodities during the 1980 - 2015 period.

A Assumptions and Regularity Conditions

In this appendix, we present the assumptions and regularity conditions that are necessary for the stable distribution characterization in Theorem 2.3. As discussed in Section 2, these assumptions admit a large class of continuous unit processes for the agents in our setup. The first assumption establishes basic integrability conditions that are common for both continuous semimartingales and Itô processes.

Assumption A.1. *For all $i = 1, \ldots, N$, the growth rate processes μ_i satisfy*

$$\int_0^T |\mu_i(t)|\, dt < \infty, \quad T > 0, \quad \text{a.s.}, \tag{A.1}$$

and the volatility processes δ_{is} satisfy

$$\int_0^T \left(\delta_{i1}^2(t) + \cdots + \delta_{iM}^2(t)\right)\, dt < \infty, \quad T > 0, \quad \text{a.s.}, \tag{A.2}$$

$$\delta_{i1}^2(t) + \cdots + \delta_{iM}^2(t) > 0, \quad t > 0, \quad \text{a.s.} \tag{A.3}$$

$$\lim_{t\to\infty} \frac{1}{t}\left(\delta_{i1}^2(t) + \cdots + \delta_{iM}^2(t)\right) \log\log t = 0, \quad \text{a.s.}, \tag{A.4}$$

Conditions (A.1) and (A.2) are standard in the definition of an Itô process, while condition (A.3) ensures that agents' holdings of units contain a nonzero random component at all times. Condition (A.4) is similar to a boundedness condition in that it ensures that the variance of agents' unit holdings does not diverge to infinity too rapidly.

The second assumption underlying our results establishes that no two agents' unit holdings be perfectly correlated over time. In other words, there must always be some idiosyncratic component to each agent's unit dynamics. Finally, we also assume that no agent's unit holdings relative to the total units for all agents shall disappear too rapidly.

Assumption A.2. *The symmetric matrix $\rho(t)$, given by $\rho(t) = (\rho_{ij}(t))$, where $1 \le i, j \le N$, is nonsingular for all $t > 0$, a.s.*

Assumption A.3. *For all $i = 1, \ldots, N$, the unit share processes θ_i satisfy*

$$\lim_{t\to\infty} \frac{1}{t} \log \theta_i(t) = 0, \quad \text{a.s.} \tag{A.5}$$

B Proofs

This appendix presents the proofs of Lemmas 2.1, 2.2 3.1, and 3.2, and Theorem 2.3.

Proof of Lemma 2.1. By definition, $x(t) = x_1(t) + \cdots + x_N(t)$ and for all $i = 1, \ldots, N$, $\theta_i(t) = x_i(t)/x(t)$. This implies that

$$dx(t) = \sum_{i=1}^N dx_i(t) = \sum_{i=1}^N \theta_i(t) x(t) \frac{dx_i(t)}{x_i(t)},$$

from which it follows that

$$\frac{dx(t)}{x(t)} = \sum_{i=1}^N \theta_i(t) \frac{dx_i(t)}{x_i(t)}. \tag{B.1}$$

We wish to show that the process satisfying equation (2.3) also satisfies equation (B.1).

If we apply Itô's Lemma to the exponential function, then equation (2.3) yields

$$\begin{aligned} dx(t) = x(t)\mu(t)\,dt + \frac{1}{2}x(t)\sum_{i,j=1}^{N}\theta_i(t)\theta_j(t)\left(\sum_{s=1}^{M}\delta_{is}(t)\delta_{js}(t)\right)\,dt \\ + x(t)\sum_{i=1}^{N}\sum_{s=1}^{M}\theta_i(t)\delta_{is}(t)\,dB_s(t), \end{aligned} \tag{B.2}$$

a.s., where $\mu(t)$ is given by equation (2.5). Using the definition of $\rho_{ij}(t)$ from equation (2.2), we can simplify equation (B.1) and write

$$\frac{dx(t)}{x(t)} = \left(\mu(t) + \frac{1}{2}\sum_{i,j=1}^{N}\theta_i(t)\theta_j(t)\rho_{ij}(t)\right)\,dt + \sum_{i=1}^{N}\sum_{s=1}^{M}\theta_i(t)\delta_{is}(t)\,dB_s(t). \tag{B.3}$$

Similarly, the definition of $\mu(t)$ from equation (2.5) allows us to further simplify equation (B.3) and write

$$\begin{aligned} \frac{dx(t)}{x(t)} &= \left(\sum_{i=1}^{N}\theta_i(t)\mu_i(t) + \frac{1}{2}\sum_{i=1}^{N}\theta_i(t)\rho_{ii}(t)\right)\,dt + \sum_{i=1}^{N}\sum_{s=1}^{M}\theta_i(t)\delta_{is}(t)\,dB_s(t) \\ &= \sum_{i=1}^{N}\theta_i(t)\left(\mu_i(t) + \frac{1}{2}\rho_{ii}(t)\right)\,dt + \sum_{i=1}^{N}\sum_{s=1}^{M}\theta_i(t)\delta_{is}(t)\,dB_s(t). \end{aligned} \tag{B.4}$$

If we again apply Itô's Lemma to the exponential function, then equation (2.1) yields, a.s., for all $i = 1, \ldots, N$,

$$\begin{aligned} dx_i(t) &= x_i(t)\left(\mu_i(t) + \frac{1}{2}\sum_{s=1}^{M}\delta_{is}^2(t)\right)\,dt + x_i(t)\sum_{s=1}^{M}\delta_{is}(t)\,dB_s(t) \\ &= x_i(t)\left(\mu_i(t) + \frac{1}{2}\rho_{ii}(t)\right)\,dt + x_i(t)\sum_{s=1}^{M}\delta_{is}(t)\,dB_s(t). \end{aligned} \tag{B.5}$$

Substituting equation (B.5) into equation (B.4) then yields

$$\frac{dx(t)}{x(t)} = \sum_{i=1}^{N}\theta_i(t)\frac{dx_i(t)}{x_i(t)},$$

which completes the proof. □

Proof of Lemma 2.2. Agents' unit holding processes x_i are absolutely continuous in the

sense that the random signed measures $\mu_i(t)\,dt$ and $\rho_{ii}(t)\,dt$ are absolutely continuous with respect to Lebesgue measure. As a consequence, we can apply Lemma 4.1.7 and Proposition 4.1.11 from Fernholz (2002), which yields equations (2.11) and (2.12). □

Proof of Lemma 3.1. This relationship between the rank-based relative growth rate parameters α_k and the local time parameters κ_k is established in the proof of Theorem 2.3 below (see equation (B.9) below). That proof also establishes the fact that $\alpha_N = -(\alpha_1 + \cdots + \alpha_{N-1})$ (see equation (B.11) below). □

Proof of Lemma 3.2. Consider the function $f_k(\theta_1, \ldots, \theta_N) = \theta_{(1)} + \cdots + \theta_{(k)}$, where $1 \leq k \leq N$. This function satisfies

$$\frac{\partial f_k}{\partial \theta_l} = 1,$$

for all $l = 1, \ldots, k$, and

$$\frac{\partial f_k}{\partial \theta_l} = 0,$$

for all $l = k+1, \ldots, N$. Furthermore, the support of the local time processes $\Lambda_{\log \theta_{(k)} - \log \theta_{(k+1)}}$ is the set $\{t \,:\, \theta_{(k)}(t) = \theta_{(k+1)}(t)\}$, for all $k = 1, \ldots, N-1$. According to Theorem 4.2.1 and equations (3.1.1)-(3.1.2) of Fernholz (2002), then, the function $f_k(\theta_1, \ldots, \theta_N) = \theta_{(1)} + \cdots + \theta_{(k)}$ satisfies the stochastic differential equation

$$\begin{aligned} d\log(x_{p_t(1)}(t) + \cdots + x_{p_t(k)}(t)) - d\log x(t) &= d\log f_k(\theta_1(t), \ldots, \theta_N(t)) \\ &\quad - \frac{\theta_{(k)}(t)}{2(\theta_{(1)}(t) + \cdots + \theta_{(k)}(t))}\, d\Lambda_{\log \theta_{(k)} - \log \theta_{(k+1)}}, \quad \text{a.s.,} \end{aligned} \tag{B.6}$$

for all $k = 1, \ldots, N$.[14] Equation (B.6) is equivalent to

$$\begin{aligned} d\log\left(\theta_{p_t(1)}(t) + \cdots + \theta_{p_t(k)}(t)\right) &= d\log\left(\theta_{(1)}(t) + \cdots + \theta_{(k)}(t)\right) \\ &\quad - \frac{\theta_{(k)}(t)}{2(\theta_{(1)}(t) + \cdots + \theta_{(k)}(t))} d\Lambda_{\log \theta_{(k)} - \log \theta_{(k+1)}}(t), \end{aligned}$$

which confirms equation (3.3) from Lemma 3.2. □

Proof of Theorem 2.3. This proof follows arguments from Chapter 5 of Fernholz (2002).

[14] Equation (B.6) relies on the fact that $\log(x_{p_t(1)}(t) + \cdots + x_{p_t(k)}(t))$ is the value over time of a "portfolio" of unit holdings with weights of $\frac{\theta_{(l)}(t)}{\theta_{(1)} + \cdots + \theta_{(k)}}$ placed on each ranked unit holding $l = 1, \ldots, k$ and weights of zero placed on each ranked unit holding $l = k+1, \ldots, N$.

According to equation (2.14), for all $k = 1, \ldots, N$,

$$\begin{aligned}\log \theta_{(k)}(T) = \int_0^T \left(\mu_{p_t(k)}(t) - \mu(t)\right)\, dt + \frac{1}{2}\Lambda_{\log\theta_{(k)} - \log\theta_{(k+1)}}(T) - \frac{1}{2}\Lambda_{\log\theta_{(k-1)} - \log\theta_{(k)}}(T) \\ + \sum_{s=1}^{M} \int_0^T \delta_{p_t(k)s}(t)\, dB_s(t) - \sum_{i=1}^{N}\sum_{s=1}^{M} \int_0^T \theta_i(t)\delta_{is}(t)\, dB_s(t).\end{aligned} \tag{B.7}$$

Consider the asymptotic behavior of the process $\log \theta_{(k)}$. Assuming that the limits from equation (2.18) exist, then according to the definition of α_k from equation (2.16), the asymptotic behavior of $\log \theta_{(k)}$ satisfies

$$\begin{aligned}\lim_{T\to\infty} \frac{1}{T} \log \theta_{(k)}(T) = \alpha_k + \frac{1}{2}\kappa_k - \frac{1}{2}\kappa_{k-1} + \lim_{T\to\infty} \frac{1}{T} \sum_{s=1}^{M} \int_0^T \delta_{p_t(k)s}(t)\, dB_s(t) \\ - \lim_{T\to\infty} \frac{1}{T} \sum_{i=1}^{N}\sum_{s=1}^{M} \int_0^T \theta_i(t)\delta_{is}(t)\, dB_s(t), \quad \text{a.s.}\end{aligned} \tag{B.8}$$

Assumption A.3 ensures that the term on the left-hand side of equation (B.8) is equal to zero, while Assumption A.1 ensures that the last two terms of the right-hand side of this equation are equal to zero as well (see Lemma 1.3.2 from Fernholz, 2002). If we simplify equation (B.8), then, we have that

$$\alpha_k = \frac{1}{2}\kappa_{k-1} - \frac{1}{2}\kappa_k, \tag{B.9}$$

which implies that

$$\alpha_k - \alpha_{k+1} = \frac{1}{2}\kappa_{k-1} - \kappa_k + \frac{1}{2}\kappa_{k+1}, \tag{B.10}$$

for all $k = 1, \ldots, N-1$. Since equation (B.9) is valid for all $k = 1, \ldots, N$, this establishes a system of equations that we can solve for κ_k. Doing this yields the equality

$$\kappa_k = -2(\alpha_1 + \cdots + \alpha_k), \tag{B.11}$$

for all $k = 1, \ldots, N$. Note that asymptotic stability ensures that $\alpha_1 + \cdots + \alpha_k < 0$ for all $k = 1, \ldots, N$, while the fact that $\alpha_N = \frac{1}{2}\kappa_{N-1} = -(\alpha_1 + \cdots + \alpha_{N-1})$ ensures that $\alpha_1 + \cdots + \alpha_N = 0$. Furthermore, if $\alpha_1 + \cdots + \alpha_k > 0$ for some $1 \leq k < N$, then equation (B.11) generates a contradiction since $\kappa_k \geq 0$ by definition. In this case, it must be that Assumption A.3 is violated and $\lim_{T\to\infty} \frac{1}{T} \log \theta_{(k)}(T) \neq 0$ for some $1 \leq k \leq N$.

The last term on the right-hand side of equation (2.15) is an absolutely continuous martingale, and hence can be represented as a stochastic integral with respect to Brownian motion

$B(t)$.[15] This fact, together with equation (3.2) and the definitions of α_k and σ_k from equations (2.16)-(2.17), motivates our use of the stable version of the process $\log\theta_{(k)} - \log\theta_{(k+1)}$. Recall that, by equation (2.19), this stable version is given by

$$d\left(\log\hat{\theta}_{(k)}(t) - \log\hat{\theta}_{(k+1)}(t)\right) = -\kappa_k\,dt + d\Lambda_{\log\hat{\theta}_{(k)}-\log\hat{\theta}_{(k+1)}}(t) + \sigma_k\,dB(t), \tag{B.12}$$

for all $k = 1, \ldots, N-1$. According to Fernholz (2002), Lemma 5.2.1, for all $k = 1, \ldots, N-1$, the time-averaged limit of this stable version satisfies

$$\lim_{T\to\infty}\frac{1}{T}\int_0^T \left(\log\hat{\theta}_{(k)}(t) - \log\hat{\theta}_{(k+1)}(t)\right)\,dt = \frac{\sigma_k^2}{2\kappa_k} = \frac{\sigma_k^2}{-4(\alpha_1+\cdots+\alpha_k)}, \tag{B.13}$$

a.s., where the last equality follows from equation (B.11).

As shown by Banner et al. (2005), the processes $\log\hat{\theta}_{(k)} - \log\hat{\theta}_{(k+1)}$ are stationary if the condition $\alpha_1 + \cdots + \alpha_k < 0$ holds, for all $k = 1, \ldots, N$. Thus, by ergodicity, equation (2.20) follows from equation (B.13). To the extent that the stable version of $\log\theta_{(k)} - \log\theta_{(k+1)}$ from equation (B.12) approximates the true version of this process from equation (2.15), then, the expected value of the true process $\log\theta_{(k)} - \log\theta_{(k+1)}$ will be approximated by $-\sigma_k^2/4(\alpha_1 + \cdots + \alpha_k)$, for all $k = 1, \ldots, N-1$. □

[15] This is a standard result for continuous-time stochastic processes (Karatzas and Shreve, 1991; Nielsen, 1999).

References

Acharya, V. V. and L. H. Pedersen (2005, August). Asset pricing with liquidity risk. *Journal of Financial Economics 77*(2), 375–410.

Altonji, J. G., A. A. Smith Jr., and I. Vidangos (2013, July). Modeling earnings dynamics. *Econometrica 81*(4), 1395–1454.

Atkinson, A. B., T. Piketty, and E. Saez (2011, March). Top incomes in the long run of history. *Journal of Economic Literature 49*(1), 3–71.

Banner, A., R. Fernholz, and I. Karatzas (2005). Atlas models of equity markets. *Annals of Applied Probability 15*(4), 2296–2330.

Banz, R. W. (1981, March). The relationship between return and market value of common stocks. *Journal of Financial Economics 9*(1), 3–18.

Benhabib, J., A. Bisin, and S. Zhu (2011, January). The distribution of wealth and fiscal policy in economies with finitely lived agents. *Econometrica 79*(1), 123–157.

Champernowne, D. G. (1953, June). A model of income distribution. *Economic Journal 63*(250), 318–351.

Cochrane, J. H. (2005). *Asset Pricing* (Revised ed.). Princeton, NJ: Princeton University Press.

Duffie, D. (2001). *Dynamic Asset Pricing Theory*. Princeton, NJ: Princeton University Press.

Fama, E. F. and K. R. French (1993, February). Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics 33*(1), 3–56.

Fernholz, E. R. (2002). *Stochastic Portfolio Theory*. New York, NY: Springer-Verlag.

Fernholz, R. T. (2016a, January). A model of economic mobility and the distribution of wealth. *mimeo, Claremont McKenna College*.

Fernholz, R. T. (2016b, January). A statistical model of inequality. *arXiv:1601.04093v1 [q-fin.EC]*.

Fernholz, R. T. and R. Fernholz (2014, July). Instability and concentration in the distribution of wealth. *Journal of Economic Dynamics and Control 44*, 251–269.

Fernholz, R. T. and C. Koch (2016, February). Why are big banks getting bigger? *Federal Reserve Bank of Dallas Working Paper 1604*.

Gabaix, X. (1999, August). Zipf's law for cities: An explanation. *Quarterly Journal of Economics 114*(3), 739–767.

Gabaix, X. (2009, 05). Power laws in economics and finance. *Annual Review of Economics 1*(1), 255–294.

Guvenen, F. (2009, January). An empirical investigation of labor income processes. *Review of Economic Dynamics 12*(1), 58–79.

Ichiba, T., V. Papathanakos, A. Banner, I. Karatzas, and R. Fernholz (2011). Hybrid atlas models. *Annals of Applied Probability 21*(2), 609–644.

Janicki, H. and E. S. Prescott (2006). Changes in the size distribution of us banks: 1960-2005. *FRB Richmond Economic Quarterly 92*(4), 291–316.

Karatzas, I. and S. E. Shreve (1991). *Brownian Motion and Stochastic Calculus*. New York, NY: Springer-Verlag.

Karatzas, I. and S. E. Shreve (1998). *Methods of Mathematical Finance*. New York, NY: Springer-Verlag.

Lucas, Jr., R. E. (1978, November). Asset prices in an exchange economy. *Econometrica 46*(6), 1429–1445.

Luttmer, E. G. J. (2007, August). Selection, growth, and the size distribution of firms. *Quarterly Journal of Economicsr 122*(3), 1103–1144.

Neumark, D., B. Wall, and J. Zhang (2011, February). Do small businesses create more jobs? new evidence for the united states from the national establishment time series. *Review of Economics and Statistics 93*(1), 16–29.

Newman, M. E. J. (2006, May). Power laws, pareto distributions, and zipf's law. *arXiv:cond-mat/0412004v3 [cond-mat.stat-mech]*.

Nielsen, L. T. (1999). *Pricing and Hedging of Derivative Securities*. New York, NY: Oxford University Press.

Pal, S. and J. Pitman (2008). One-dimensional brownian particle systems with rank-dependent drifts. *Annals of Applied Probability 18*(6), 2179–2207.

Piketty, T. (2014). *Capital in the Twenty-First Century*. Cambridge, MA: Harvard University Press.

Saez, E. and G. Zucman (2014, October). Wealth inequality in the United States since 1913: Evidence from capitalized income tax data. *NBER Working Paper 20625*.

Shkolnikov, M. (2011). Competing particle systems evolving by interacting lévy processes. *Annals of Applied Probability 21*(5), 1911–1932.

Van Nieuwerburgh, S. and P.-O. Weill (2010, October). Why has house price dispersion gone up? *Review of Economic Studies 77*(4), 1567–1606.

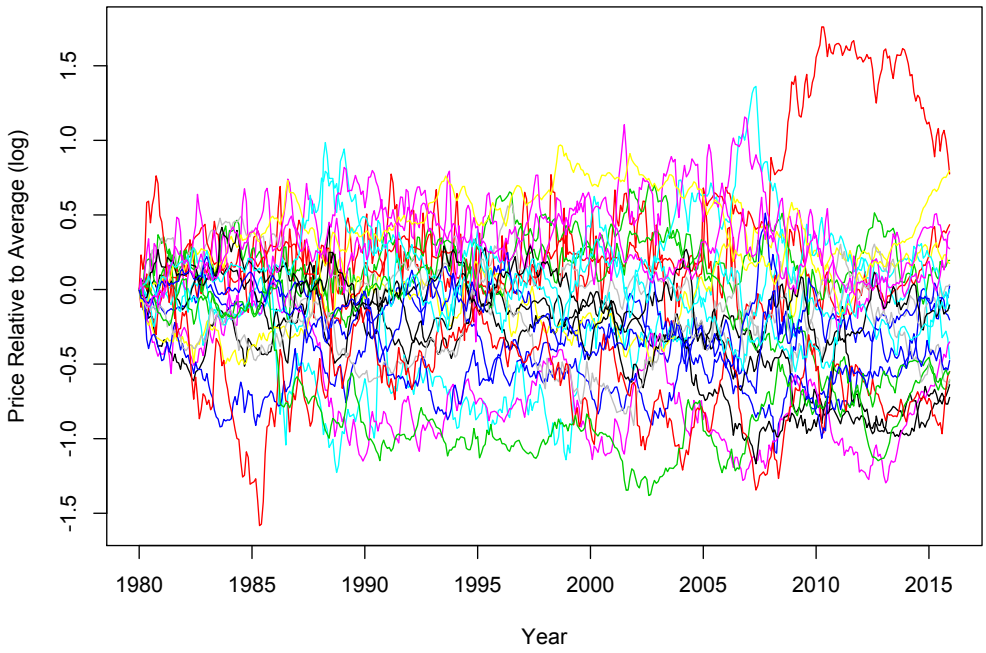

Figure 1: Log prices of commodities relative to the average price of all commodities, 1980 - 2015.

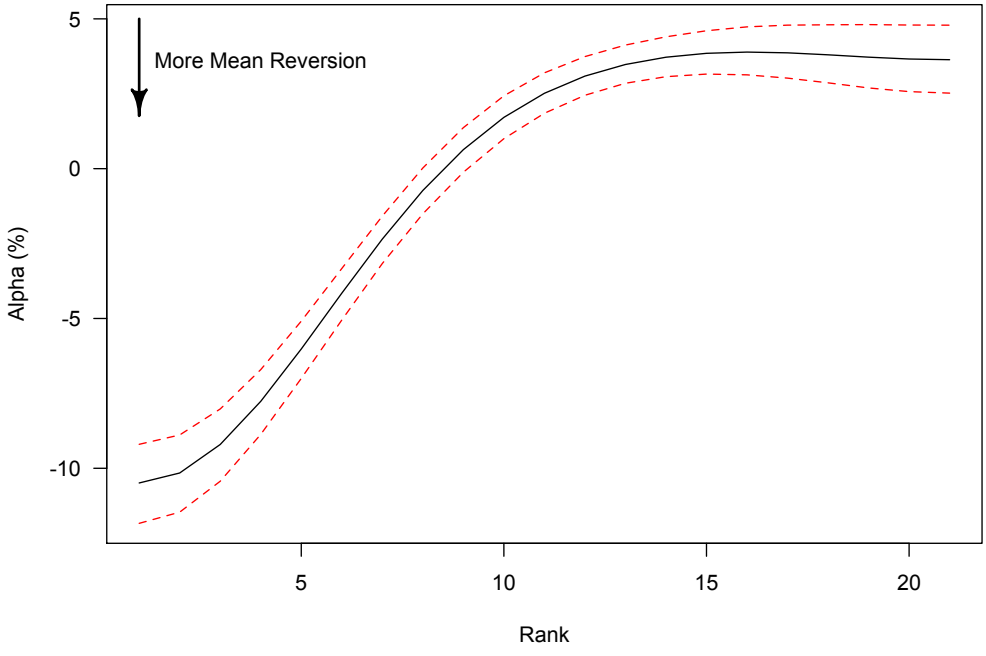

Figure 2: Point estimates and 95% confidence intervals of minus the reversion rates (α_k) for different ranked commodities, 1980 - 2015.

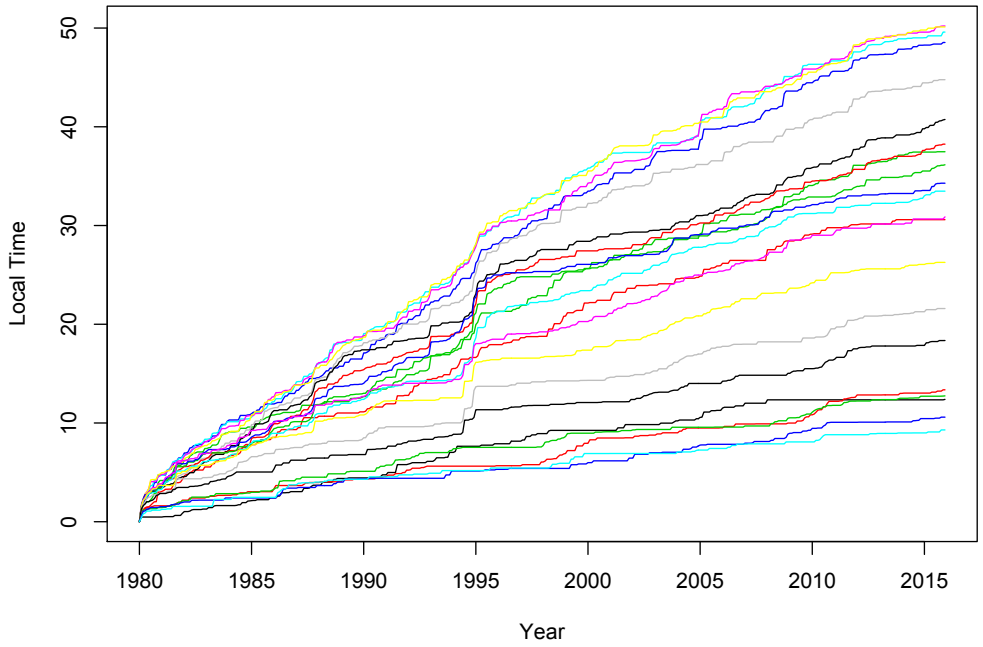

Figure 3: Local time processes ($\Lambda_{\log \tilde{x}_{(k)} - \log \tilde{x}_{(k+1)}}$) for different ranked commodities, 1980 - 2015.

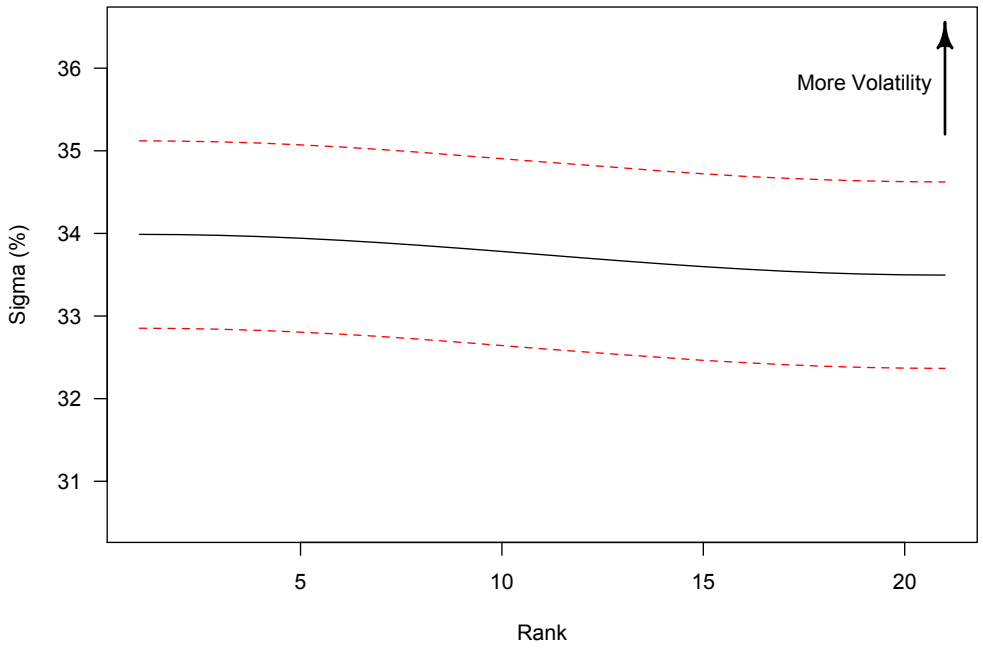

Figure 4: Point estimates and 95% confidence intervals of standard deviations of idiosyncratic commodity price volatilities (σ_k) for different ranked commodities, 1980 - 2015.

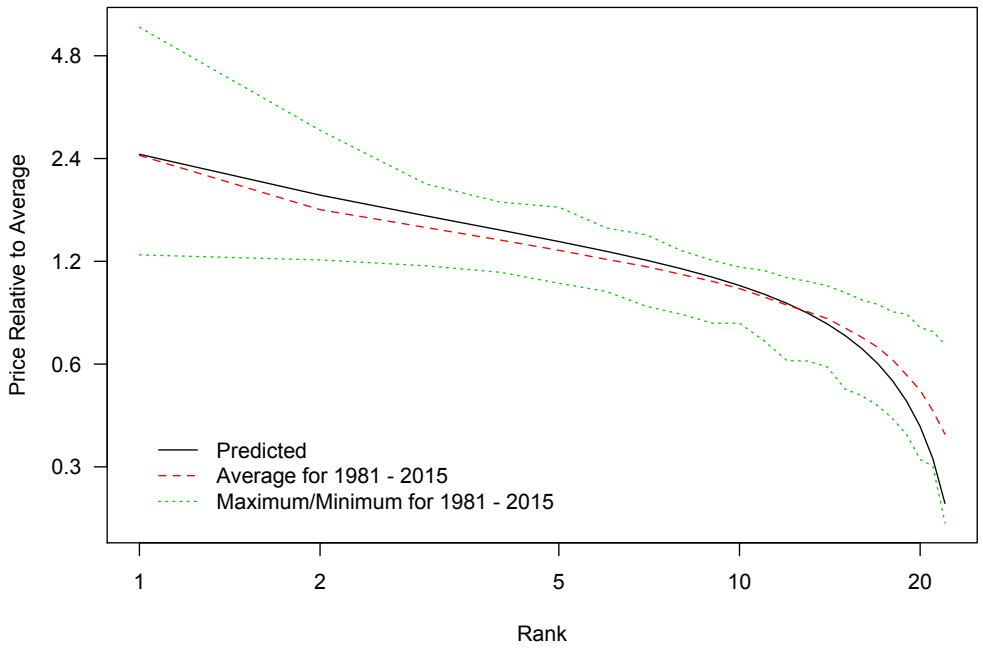

Figure 5: Relative commodity prices for different ranked commodities for 1981 - 2015 as compared to the predicted relative prices.

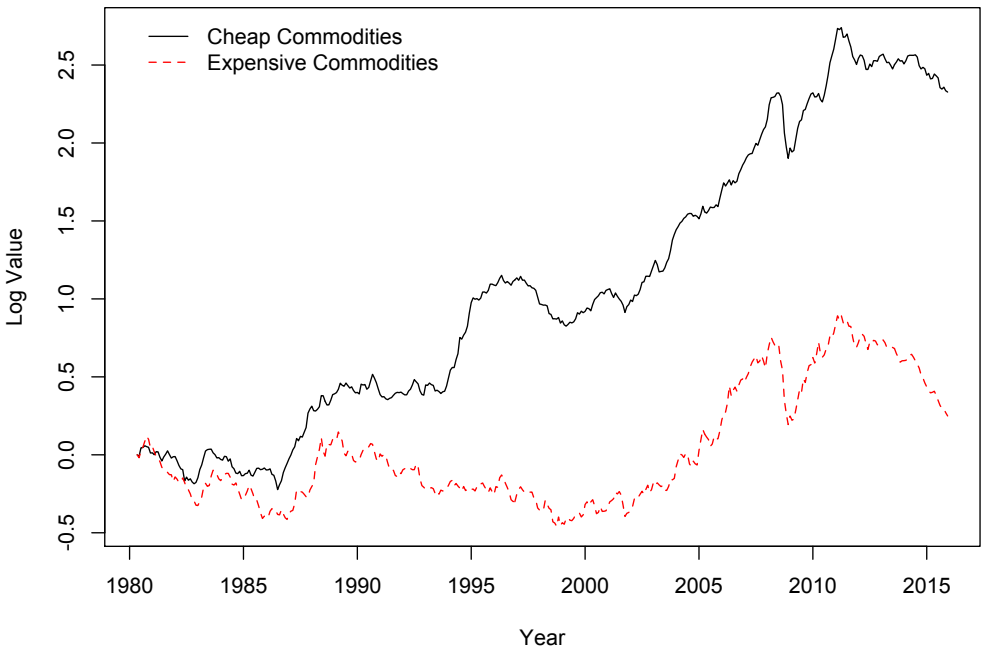

Figure 6: Log returns for cheap-commodities and expensive-commodities portfolios, 1980 - 2015.

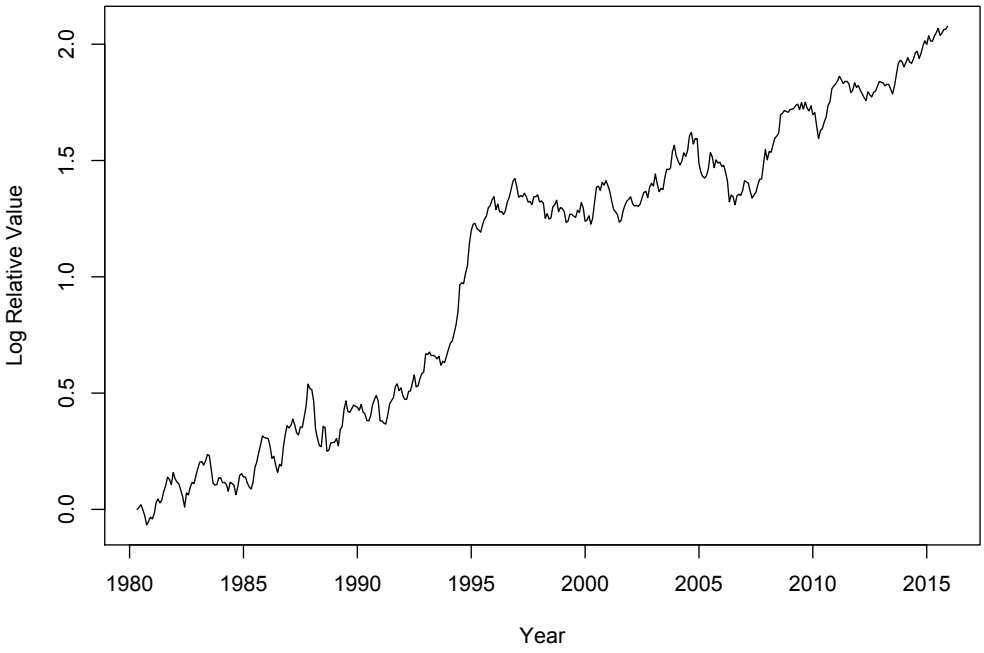

Figure 7: Log return of cheap-commodities portfolio relative to expensive-commodities portfolio, 1980 - 2015.